

CLAIMS

1. A method for dynamically adjusting the workload of an active resource, the workload being expressed as a collection of units, each unit including its own key identifier, the active resource being associated with at least one parent workload group, the parent workload group including a collection of workload units such that workload units belonging to the parent workload group share an identical sequence of values at a specified depth value of their key identifiers, the identical sequence of values defining a group key identifier associated with the parent workload group, the method comprising:
 - 5 independently determining by the active resource that an overload condition exists at the active resource; and
 - if the overload condition exists:
 - increasing the depth value of the parent workload group such that at least two child workload groups are identified; and
 - 15 assigning a target resource to manage at least one of the child workload groups.
2. The method of claim 1, further comprising if the overload condition exists, identifying at least one candidate resource to which the child workload groups may be distributed.
3. The method of claim 1, further comprising requesting workload acceptance from the target resource.
4. The method of claim 1, further comprising recording the parent workload group as inactive at the active resource.
5. The method of claim 1, further comprising transferring application-specific objects corresponding to the child workload groups.
6. The method of claim 1, further comprising redirecting entities operating on elements of the parent workload group to the target resource managing the child workload group.
7. The method of claim 6, further comprising:

receiving a probe message from an entity operating on a workload unit that is a member of the parent workload group, the probe message including a guessed identifier key formed by guessing a depth to be
5 associated with the unit's key identifier; and

sending a response to the entity indicating the group key identifier that the current resource locally determines to be the nearest known active parent group to which the element's key identifier belongs.

8. The method of claim 7, wherein the entity operating on a workload unit uses the response to further refine its estimate of a correct depth to be associated with the unit's key identifier; and
probing another resource associated with the parent key group
5 formed by using the refined depth of the unit's key identifier.

9. The method of claim 1, further comprising:
determining that an under-load condition exists at the active resource; and

if the under-load condition exists:
5 identifying at least two workload groups for consolidation into a consolidated workload group;

generating a consolidated key identifier such that workload units belonging to the consolidated workload group share an identical sequence of values at a specified depth value of the consolidated key
10 identifier; and

managing the consolidated workload group by the active resource.

10. The method of claim 9, wherein generating the consolidated key identifier includes decreasing the depth value of the parent workload group such that the consolidated workload group is identified.

11. The method of claim 1, further comprising associating the workload unit with the key identifier such that the key identifier encodes one or more attributes of the workload unit.

12. The method of claim 1, further comprising constructing a virtual key for mapping to the target resource, wherein the virtual key

includes a load-dependent subset of a complete key identifier.

13. The method of claim 12, further comprising using the constructed load-dependent virtual key as an input to a separate mapping service that returns the identity of the target resource to which the workload elements belonging to the virtual key should be
5 directed.

14. A system for running a distributed computer application whose workload can be decomposed into a set of workload units, each workload unit including its own key identifier, over a dynamically varying set of distributed resources, the number of resources involved
5 in the distributed computation varying dynamically in response to changes in an overall workload, the system comprising:

a set of active resources cooperatively managing an entire set of key identifiers constituting the overall workload, each individual resource managing a dynamically varying group of key identifiers, each
10 resource independently evaluating its own workload condition and deciding on the creation or consolidation of group key identifiers to reduce or increase its workload;

an overall set of resources, of which the active resources constitute a subset that can be utilized as part of the distributed
15 computer application as needed;

a mapping service configured to receive a virtual key associated with at least one key identifier as input and configured to produce the identity of the target resource from the overall resource set as an output; and

20 a set of client entities utilizing the distributed application, each client entity being associated with at least one key identifier, and each client entity dynamically determining the load-dependent group of identifier keys that each such currently belongs to.

15. A system for dynamically adjusting the workload of an active resource, the system comprising:

a plurality of workload units, each workload unit being

associated with a key identifier;

5 a parent workload group including a collection of workload units, wherein the workload units belonging to the parent workload group share an identical sequence of values at a specified depth value of the key identifier; and

10 an active resource managing the parent workload, the active resource configured to increase the depth value of the parent workload group such that at least two child workload groups are identified and assign a target resource to manage at least one of the child workload groups if an overload condition exists at the active resource.

16. The system of claim 15, wherein the active resource is further configured to identify at least one candidate resource to which the child workload groups may be distributed.

17. The system of claim 15, wherein the active resource is further configured to request workload acceptance from the target resource.

18. The system of claim 15, wherein the active resource is further configured to record the parent workload group as inactive at the active resource.

19. The system of claim 15, wherein the active resource is further configured to transfer application-specific objects corresponding to the child workload groups.

20. The system of claim 15, wherein the active resource is further configured to identify at least two workload groups for consolidation into a consolidated workload group, generate a consolidated key identifier such that workload units belonging to the
5 consolidated workload group share an identical sequence of values at a specified depth value of the consolidated key identifier, and manage the consolidated workload group by the active resource if an under-load condition exists.

21. The system of claim 20, wherein the active resource is

further configured to decrease the depth value of the parent workload group key identifier such that the consolidated workload group is identified.

22. The system of claim 15, wherein the active resource is further configured to associate the workload unit with the key identifier such that the key identifier encodes one or more attributes of the workload unit.

23. The system of claim 15, wherein the active resource is further configured to construct a virtual key for mapping to the target resource, wherein the virtual key includes a load-dependent subset of a complete key identifier .

24. The system of claim 15, further comprising an external service configured to identify at least one candidate resource to which the child workload groups may be distributed.

25. A computer program product embodied in a tangible media comprising:

computer readable program codes coupled to the tangible media for dynamically adjusting the workload of an active resource, the active
5 resource being associated with at least one parent workload group, the parent workload group including a collection of workload units, each workload unit being associated with a key identifier such that workload units belonging to the parent workload group share an identical sequence of values at a specified depth value of the key identifier,
10 the identical sequence of values defining a group key identifier associated with the parent workload group, the computer readable program codes configured to cause the program to:

determine that an overload condition exists at the active resource; and
15 if the overload condition exists:
increase the depth value of the parent workload group such that at least two child workload groups are identified; and
assign a target resource to manage at least one of the child

workload groups.

26. The computer program product of claim 25, further comprising program code configured to identify at least one candidate resource to which the child workload groups may be distributed if the overload condition exists.

27. The computer program product of claim 25, further comprising program code configured to request workload acceptance from the target resource.

28. The computer program product of claim 25, further comprising program code configured to record the parent workload group as inactive at the active resource.

29. The computer program product of claim 25, further comprising program code configured to transfer application-specific objects corresponding to the child workload groups.

30. The computer program product of claim 25, further comprising program code configured to redirect entities operating on elements of the parent workload group to the target resource managing the child workload group.

31. The computer program product of claim 30, further comprising program code configured to:

- 5 receive a probe message from an entity operating on elements of the parent workload group, the probe message including a guessed identifier key at a guessed depth; and
- send a response to the entity indicating the target resource managing the child workload group.

32. The computer program product of claim 25, further comprising program code configured to:

- determine that an under-load condition exists at the active resource; and
- 5 if the under-load condition exists:
 - identify at least two workload groups for consolidation into a

consolidated workload group;

generate a consolidated key identifier such that workload units
belonging to the consolidated workload group share an identical
10 sequence of values at a specified depth value of the consolidated key
identifier; and

manage the consolidated workload group by the active resource.

33. The computer program product of claim 31, wherein the
program code to generate the consolidated key includes program code to
decrease the depth value of the parent workload group such that the
consolidated workload group is identified.

34. The computer program product of claim 25, further
comprising program code configured to associate the workload unit with
the key identifier such that the key identifier encodes one or more
attributes of the workload unit.

35. The computer program product of claim 25, further
comprising program code configured to construct a virtual key for
mapping to the target resource, wherein the virtual key includes a hash
value of the key identifier.